

METHOD AND APPARATUS FOR MANAGING REMOTE IP
NETWORK ELEMENTS THROUGH SONET NETWORK ELEMENTS

BACKGROUND OF THE INVENTION

Networks of various companies and agencies have the potential to become integrated with a standardized fiber optic plant to promote efficient sharing and transmission of information. However, the array of different protocols used by different bodies hinders interconnectivity and management of data through a common fiber facility.

There is a current need for local and wide area telecommunications networks to connect Internet Protocol (IP) managed Network Elements (NEs) at customer premises with Synchronous Optical NETWORK elements (SONET NEs), which are external to the customer premises network and provide a data transport mechanism. As shown in Fig. 1, a SONET NE network including OSI-DCC has the potential to transport data to and from the customer premises where IP NEs 5₁, 5₂ and 5₃ are used to provide the customer a specific interface or service (ISDN, HDLC). These IP NEs (5₁, 5₂ and 5₃) may need to be managed from an IP manager which is remote from the local customer premises. However, the interconnectivity of remotely located IP NEs, such as those represented by reference numerals 5₁ through 5₃, through an intermediate SONET NE 3 to the IP management device 1 has not been previously realized due to fundamental differences in the IP and OSI (used by SONET) protocol stacks.

For instance, the SONET NEs are managed from a central location via Open Systems Interconnect (OSI) applications running on an overhead management channel called the data channel communications (DCC). The customer premises NEs may, however, be IP NEs that are managed by IP applications such as a Simple Network Management Protocol (SNMP). Presently, there is no generic way to manage these remote IP NEs using SNMP through the existing SONET OSI-DCC.

Additionally, at the network layer, SONET runs the CLNP protocol while the IP NE runs the TCP/IP protocol. When a conventional SONET device encounters an IP packet, the SONET device CLNP layer will not be able to route the packet to the IP destination address.

DESCRIPTION OF THE RELATED ART

5 A first conventional technique to permit interconnection between the customer premise IP NE 5₁-5₃ and a remote IP management device 1 can be provided using a separate network for carrying this management information, using a modem and a dedicated line. This is illustrated in Figure 2, where the IP management device 1 running SNMP is connected via a leased line 3 to the IP NEs 5₁-5₃ to be managed via modem connections 2, 4.

10 An obvious disadvantage of this first conventional technique is that it requires the use of a dedicated leased line (DSO) and the addition of three extra pieces of equipment (two modems and a terminal server) that must be managed and maintained.

In a second conventional technique to permit management of remotely located IP NEs, management information is embedded in a data path DS1 of one of the remote IP
15 devices. This is illustrated in Fig. 3. For each group of n remote IP NEs (5₁ to 5₃) to be managed, a similar IP NE 2 would be required in the OSI-SONET management office and connected to the management station 1. In this scenario, the IP NEs would provide their own proprietary management channel in band to communicate between the local IP NEs (directly connected the management station) and the remote IP NEs. The second method would
20 require the purchase of additional equipment that must be placed in a central management office where space is a premium, and also reduces the available bandwidth that can be sold to an end user customer.

SUMMARY OF THE INVENTION

The present invention obviates the above deficiencies of the known techniques by
25 tunneling management data and other information to and from remotely located IP NEs via

the SONET DCC by placing IP over the Connection Less Network Protocol (CLNP) that is present in OSI.

The placement of IP over CLNP is generic enough to permit pass through management of any IP protocol and is not specific to any vendor. Moreover, because this protocol interfaces at CLNP, IP over CLNP will pass unhindered though the OSI stack on the DCC of legacy SONET NEs or SONET NEs of other vendors that do not have the modification to transmit IP packet information in addition to CLNP datagrams. Thus, the OSI network is not effected by the overlaid IP network. While a few benefits have been described above, additional advantages of the present invention will be apparent to one skilled in the art in view of the following description of a preferred embodiment of the invention with reference to the drawings.

DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates remotely located IP devices with an OSI-DCC network disposed therebetween;

Fig. 2 illustrates a first conventional arrangement for interconnecting remote IP NEs;

Fig. 3 illustrates a second conventional arrangement for interconnecting remote IP NEs;

Fig. 4 illustrates the conceptual interconnection between remotely located IP NEs via an OSI NE according to a preferred embodiment of the invention;

Fig. 5 illustrates the peer to peer communications between layers of the IP gateways and an intermediate SONET NE according to the preferred embodiment;

Fig. 6 illustrates the interaction between an IP layer overlaid on a CLNP layer according to the preferred embodiment of the invention.

DESCRIPTION OF PREFERRED EMBODIMENT

Referring to Fig. 4, one or more SONET NEs that are used as routing nodes between remotely located IP NE devices will operate as IP gateways. More particularly, an IP management device IP₁ uses the SONET NEs (with OSI DCC) intermediate network elements, designated by NSAP₁ and NSAP₃ as routing nodes to communicate with remotely located IP network elements IP₂-IP₄ to be managed. Similarly, in order to route IP information from the device IP₁ to IP network elements IP₅, IP₆, the OSI DCC network elements designated by NSAP₁ and NSAP₅ are used as routing nodes. In Fig. 4, the OSI DCC elements NSAP₁, NSAP₃ and NSAP₅ act as IP gateways and include an IP stack with additional software code to "tunnel" between IPGs. The IP gateway stack permits the OSI DCC device to receive a conventionally formatted IP packet and route it through an IP tunnel over OSI on the DCC to another OSI DCC device which is directly connected with the destination IP device. The details of how the IP routing tunnels become implemented will be described in greater detail below, in conjunction with Fig. 6.

Referring first to Fig. 5, an IP gateway (such as IPG₁, of Fig. 4) receives IP information packets from an IP device via known TCP/IP protocols. The Network Access Protocol (NAP) layer extracts its header and trailer information and passes the IP information to the upper IP layer. These interactions occur according to known interoperations and no additional description is provided. The IP information, including an IP destination address, is used in the network interface (NI) layer according to the present invention. The use of this information permits tunneling to the specific OSI NE which is attached to the IP device with the requested destination address. For purposes of this disclosure, the IPG that initially receives the IP packet will be referred to as the local IP gateway (LIG), and the IPG that delivers the IP packet to the destination IP address will be referred to as the remote IP gateway (RIG).

The LIG achieves tunneling communication with the OSI device RIG that is attached to the destination IP device. The LIG to RIG communication is OSI and the present invention makes use of layer 3 (CLNP) of the seven layer OSI stack. Once the LIG to RIG communication is achieved, the IP stack in the RIG and LIG view this link as just another routable link that is managed by standard IP routing protocols such as OSPF or RIP. This completes the communication between the IP end devices using the OSI devices as intermediate IP gateways. The intermediate NE may include any SONET NE that has OSI running on its DCC.

Notably, the communication at the CLNP layer permits regular operation of any SONET device with OSI on the DCC even if they do not have this tunneling feature. Therefore, the SONET network is transparent to the IP over CLNP protocol implemented by the present invention. Those OSI DCC devices which are called upon to act as IPGs also will make use of a selector field (0xf0) to distinguish the handling of incoming IP and traditional OSI traffic. If an OSI packet is received, then the packet is delivered to the transport layer of the OSI protocol. If an IP packet is received, then the packet is routed to the IP tunneling interface according to the present invention.

Fig. 6 illustrates the interoperations between the IP gateway protocol stack operating as an LIG and the corresponding OSI CLNP layer. In the preferred embodiment, the IP stack of the internet gateways run an open shortest path first routing protocol (OSPF) to configure the network. This permits each IPG to ascertain knowledge about the IP addresses of the IP devices which are directly connected to the respective IPG. The OSPF configuration of the IP devices is done in a conventional manner. In relevant part, the OSPF software associates the IP addresses with a port number. At a bottom layer, a line driver element NILan interfaces with the IP LAN network. (The NiSmem is used for internal communications within the NE.) An internet protocol tunneling layer network interface (IPT_NI) interfaces

the IP layer with the CLNP layer in the LIG, which then communicates with a corresponding RIG as will be described in further detail below. Though Fig. 6 shows several IPT_NI tunnels but a connection at only one CLNS interface, each IPT_NI actually tunnels to different SONET OSI devices acting as an RIG. The IP layer may be implemented by a known protocol stack, such as PNA, produced by Integrated Systems, Inc. Similarly, the OSI functions can be implemented by known software using the specific CLNP selector, as previously mentioned in this disclosure.

The network administrator creates an IP tunnel between a LIG and a RIG. This is done by notifying the LIG of the NSAP or TID of the RIGs. Once the link is created, the IP routing protocols (OSPF and/or RIP) running on the LIG and RIG exchange information that allows them to populate their IP routing tables with information about the IP devices connected to each of the other IPGs. For instance, referring back to Fig. 4, IPG₁ can facilitate IP communications between IP₁ and 1) IP₂-IP₄ attached to IPG₂ and 2) IP₅-IP₆ attached to IPG₃.

At least two tables will be needed to make this possible. The first table (Table 1, below) is used to map IP tunnel numbers to the NSAP of the NE that terminates this tunnel. The network administrator creates the tunnel, or port connection, to a particular NSAP which acts as an IPG. This functionality can be implemented in software. The second table (Table 2, below) is the IP routing table that associates an IP destination address with a specific tunnel, or a port number. This second table is created via an IP routing protocol such as OSPF or RIP. As an exemplary case, in Fig. 6, port number 3 in the IP stack is assigned to receive and transmit information to and from the IPG identified as NSAP₃

Example :

Table 1

IPT IF 1 (Port No. 3)	NSAP ₃
IPT IF 2 (Port No. 4)	NSAP ₅

5

Table 2

IP ₂ , IP ₃ , IP ₄	IPT IF 1 (Port No. 3)
IP ₅ , IP ₆	IPT IF 2 (Port No. 4)

Together, the first lines of Tables 1 and 2 direct the LIG to transmit an IP packet via IP tunnel 1 to the OSI DCC device uniquely identified by NSAP₃ when the IP packet includes one of the IP addresses for IP₂-IP₄ as a destination address. As previously indicated, the RIG and LIG communicate via CLNP. The RIG receives the IP packet and its own routing table recognizes that IP₂, IP₃ or IP₄ are directly connected to the RIG and routes the packet accordingly via TCP/IP.

Similarly, the second line of table information of Tables 1 and 2 direct the LIG to transmit an IP packet via IP tunnel 2 to the OSI DCC device identified by NSAP₅ when the IP packet includes the IP addresses for IP₅-IP₆ as a destination address. Device NSAP₅, acting as an RIG, receives the packet recognizes that IP₅ and IP₆ are directly connected to the RIG, and routes the packet accordingly.

These tunnels identified as tunnel interfaces 1 and 2 are bi-directional and therefore the OSI DCC devices identified as NSAP₃ and NSAP₅ will use tunnel interface 1 and tunnel interface 2 respectively to communicate with IPG₁. It should be noted however that the

tunnel name/number has local significance only and therefore a LIG and RIG may refer to the same logical tunnel by different names/numbers.

It is further noted that no additional tunnel would be required between IPG2 and IPG3 since they can communicate with each other through existing tunnel interfaces 1 and 2. The

5 OSI DCC devices identified by NSAP₂ and NSAP₄ are not affected by the tunneling interfaces in their operation.

A network administrator or user may provide the table entries for the tunnel interfaces of Table 1, above. Alternatively, this may be programmed as part of the IPT network interface software. An IP tunnel manager (ITM) of the LIG requests network service access
10 points NSAP's for NE's with IPG functionality. This would include intermediate SONET NE's connected to IP end devices. A broadcast request is sent to the NE's with the IP addresses to determine if the network element is operable as an IP gateway. In other words, a message is sent to the network elements to see which ones can support the IP over CNLP association. For those NEs that can support the IP gateway function, a network interface
15 manager (NIM) on the LIG creates network interfaces for each respective NSAP. This network interface corresponds to the table entries as described above.

In addition, a new IPT_NI may be created at the RIG if it receives an IP packet from CLNP and there is no corresponding IPT_NI. This case occurs when the RIG receives an IP packet embedded in the CLNP. The CLNP datagram includes its originating NSAP. If the
20 RIG's routing table does not include a tunnel entry for the received NSAP, then the RIG may assign a port number as a tunnel interface to the originating NSAP address.

While the above provides a description of the preferred embodiment of the invention, the invention is not limited thereto and can be modified by one skilled in the art to reflect the spirit and scope of the appended claims.